



Detection of severe obstructive sleep apnea through voice analysis

Jordi Solé-Casals^{a,*}, Cristian Munteanu^{b,c}, Oriol Capdevila Martín^c, Ferran Barbé^{d,e}, Carlos Queipo^g, José Amilibia^h, Joaquín Durán-Cantolla^{e,f}

^a Data and Signal Processing Group, University of Vic – Central University of Catalonia, Sagrada Família 7, 08500 Vic, Spain

^b Universidad de Las Palmas de Gran Canaria, Campus de Tafira, 35017 Las Palmas, Spain

^c Cooclea, S.L., Mestre Garriga 10, 08500 Vic, Spain

^d Respiratory Department IRBLleida, Av. Alcalde Rovira Roure, 80, 25198 Lleida, Spain

^e CIBERES, ISCIII, Sinesio Delgado 20, Madrid, Spain

^f Sleep Unit, Service of Pneumology, Hospital Txagorritxu, Servicio Vasco de Salud-Osakidetza, José Achótegui s/n, Vitoria-Gasteiz, Spain

^g Respiratory Department, Sleep Unit – Hospital Universitario Marqués de Valdecilla, Avda. Valdecilla n° 25, 39008 Santander, Spain

^h Respiratory Department, Sleep Unit – Hospital Universitario de Cruces, Plaza de Cruces, s/n 48903 Barakaldo, Bizkaia, Spain



ARTICLE INFO

Article history:

Received 19 October 2012

Received in revised form 4 April 2014

Accepted 16 June 2014

Available online 3 July 2014

Keywords:

Obstructive sleep apnea

Voice processing

Genetic algorithms

Feature reduction

ABSTRACT

This paper deals with the potential and limitations of using voice and speech processing to detect Obstructive Sleep Apnea (OSA). An extensive body of voice features has been extracted from patients who present various degrees of OSA as well as healthy controls. We analyse the utility of a reduced set of features for detecting OSA. We apply various feature selection and reduction schemes (statistical ranking, Genetic Algorithms, PCA, LDA) and compare various classifiers (Bayesian Classifiers, kNN, Support Vector Machines, neural networks, Adaboost). S-fold crossvalidation performed on 248 subjects shows that in the extreme cases (that is, 127 controls and 121 patients with severe OSA) voice alone is able to discriminate quite well between the presence and absence of OSA. However, this is not the case with mild OSA and healthy snoring patients where voice seems to play a secondary role. We found that the best classification schemes are achieved using a Genetic Algorithm for feature selection/reduction.

© 2014 Elsevier B.V. All rights reserved.

Introduction

Obstructive Sleep Apnea Hypoapnea Syndrome (OSA for short) is a common sleep disorder that manifests itself by daytime sleepiness caused by a cease in breathing occurring repeatedly during sleep, often for a minute or longer and as many as hundreds of times during a single night. OSA is associated with a reduced-calibre upper airway, and repetitive effects of apneas and hypopneas include oxygen desaturation, reductions in intrathoracic pressure, and central nervous system arousals [1]. Diagnosis of the sleep condition is based on the calculation of the apnea-hypopnea index (AHI) which measures the frequency of reductions in air-flow associated with upper-airway collapse or narrowing that occurs with the state change from wakefulness to sleep [1]. The gold standard procedure to determine the AHI is polysomnography, however it is a quite costly methodology [2]. No other measure has proven to be superior to AHI in assessing the overall effect of obstructive sleep apnea. Nevertheless, there is no common consensus between laboratories regarding its definition. Other metrics such as the number or frequency of arousals during a night sleep might be considered an equally good indicator of OSA [1]. Thus, seeking alternative methods of diagnosis that are simpler and more cost effective is fully motivated, and in recent years it was advocated that voice may play a central role into detection of OSA syndrome. Preliminary findings on speech disorder in OSA have been reported firstly in [3] employing a rather small sample (39 subjects) and subjective results of acoustic evaluation of voice changes in OSA, followed by a

study [4] on a bigger sample (252 patients) giving again only subjective judgement results. An attempt to a more objective evaluation study was given in [5]. To discriminate between OSA patients and controls, the authors apply spectral analysis to vowels, but again the sample taken into account is small (28 subjects). Recently, in [6] and [7] the authors show the importance of using voice as a discriminatory factor for detection of severe sleep apnea employing Gaussian Mixture Models on phrases (in [6]) and on vowels (in [7]). However, the authors recognize the need for a wider training and validation sets. So far, either due to small samples or subjective judgements, it is hard to quantify up to what extent or under what circumstances we might consider voice as a good discrimination measure between OSA and healthy subjects. Recent efforts such as [8] try to model the upper-airway in OSA subjects as compared to controls by employing computational fluid dynamics models, and they conclude that there is a clear tendency to closure of the upper-airway in OSA. As the upper-way coincides in part with the vocal tract, the thinning of the lumen and tendency to closure experienced in OSA do suggest that there may be an identifiable dysfunction in voice also.

Method

Subjects

We have 376 subjects that undertook this study, both controls (proven healthy subjects) as well as snoring OSA suspects, mild OSA and severe OSA patients, 123 women and 253 men, with ages comprised between 18 and 82. This cross sectional data has been

* Corresponding author. Tel.: +34 93 8815519.

E-mail address: jordi.sole@uvic.cat (J. Solé-Casals).

pooled from several state hospitals in Spain (namely from Vitoria, Lleida, Cruces and Valdecillas). The diagnosis for each patient was confirmed by specialized medical staff through polysomnography (PSG) or through respiratory polygraphy (RP) whenever PSG was not available. For the present study we consider $AHI \leq 5$ as controls (healthy subjects) and $AHI \geq 30$ as severe OSA patients, which is in agreement with the recommendations made by the American Academy of Sleep Medicine [9]. For the purpose of clarity, along the present study, we call these subjects *extreme cases*, while in-between we may have mild OSA, or snoring non-OSA patients. Thus, among the total of 376 available cases we extract a group of 127 controls and a group of 121 severe OSA with the following characteristics.

Voice database

Speech was recorded using an AKG Perception 100 condenser microphone, a Digidigging M-box[®] sound card (Avid Audio), and a sound acquisition software by Pro Tools[®] (Avid Audio). The microphone was held 20 cm away from the subject's mouth, by a technician designated for this task. The audio signal was sampled at 44.1 kHz with 16 bits per sample, and recording was done for two distinct positions for each subject: upright or seated ('A' position) and supine or stretched ('E' position). Before each recording session, during 3 min the patient was kept as comfortable as possible in order to induce a relaxation feeling as stress is known to affect voice [10]. The room's ambient was kept quiet, in dim comfortable light and no external noise. Each subject was asked to emit the 5 vowels present in Spanish language that are: /a/, /e/, /i/, /o/, /u/ in a sustained fashion for at least 4 s each. Additionally, the patients were asked to utter the following sentence (in Spanish): \De golpe nos quedamos a oscuras\. Between each utterance a silence gap of 2–3 s was enforced through the recording protocol. The reason for using two distinct uttering positions ('A' and 'E') was that as gravity and head position affect differently the vocal tract when seated and when stretched, the sound properties also change [11,12]. Therefore, we add a second source of information per patient besides the utterance in the more common position (seated). To the best of our knowledge, this is the first attempt to detect OSA through voice analysis that uses this idea. All recordings are done by technicians from the sleep units in the 4 hospitals participating in the study, all technicians being "blind" with respect to the outcome of the experiment.

Voice features

A total of 253 features per patient were extracted from the utterance of 5 vowels and a sentence in two distinct positions. The rationale behind choosing the following listed features is that most of these measures have been previously employed for detection or characterization of pathological voice. Our working hypothesis is that severe OSA may present abnormalities in the voice production, such as increased nasality, harshness or dullness, which is also in agreement with previous findings (see [3–6]). The features may be grouped as follows.

Formant and pitch based

For each vowel we compute the second formant using the classical algorithm of root finding for the Linear Predictive Coefficient polynomial [13], with a previous octave-jump filtering step. Next, we extract the Mean Frequency (MF), Coefficient of Variation in Frequency (CVF), Jitter Factor (JF), Relative Average Perturbation (RAP), Mean Bandwidth (MBW) and Coefficient of Variation of the Bandwidth (CVBW). Definitions of these measures are given for example in [14,15]. Voice pitch is extracted for each vowel employing an improved autocorrelation method given in [16].

The post-processing octave-jump filtering stage and the features extracted from pitch are exactly the same as in the case of the second formant.

Time domain analysis

The time signal (one signal for each vowel and each subject position) yields a set of features that are pitch-synchronous in that we take as a reference signal the pitch extracted in Formant and pitch based section. The features (see [17] for detailed definitions) are the Mean Intensity/Amplitude (MIA), the Coefficient of Variation of the Intensity/Amplitude, the Shimmer of the signal Intensity (SIA) and a measure of the perturbation in the signal amplitude: Amplitude Perturbation Quotient (APQ).

Voice harshness and turbulence analysis

The first measure employed is related to the content of harmonics present in voice (versus non-harmonics content, denoted as noise) and is commonly designated as Harmonics to Noise Ratio (HNR). To compute HNR we took a well-established frequency method described in [18] among other more basic variants such as [19,20]. A particularly useful feature as turned-out to be from results obtained (see Results section) is the MHNHR: the mean HNR computed at the beginning (approximately the first second) of vowel /a/. Other measures are the Soft Phonation Index (SPI) and the Voice Turbulence Index (VTI). VTI measures the turbulence components caused by incomplete or loose adduction of the vocal folds; SPI evaluates the poorness of high-frequency harmonic components that may be an indication of loosely adducted vocal folds during phonation. In our implementation we compute SPI and VTI according to definitions in [14] but employing the improved algorithm in [18] to calculate the intra-harmonic and inter-harmonic energies present in the voice signal.

Linear prediction analysis

Based on a linear predictions analysis on the voice signal, we extracted the Pitch Amplitude (PA) and Spectral Flatness Ratio (SFR) with methods described in [21]. PA measures the dominant peak of the residual signal auto-correlation function, and SFR quantifies the flatness of the residue signal spectrum.

Dynamical systems analysis

To account for significant non-linear and non-Gaussian random phenomena present in disordered sustained vowels we employ two features inspired by dynamical system analysis performed on the voice signal. These features were introduced in [22]. The authors apply state-space recurrence analysis to produce an entropy measure H_{norm} , and Fractal scaling analysis that yields a measure called Detrended Fluctuation Analysis (DFA).

LTAS based

So far, we introduced features computed on sustained vowels. Next, we present features extracted from phrase analysis. The core analysis method of the sentence was Long-Term Average Spectrum (LTAS). In [23–25] the authors focus on the use of LTAS to quantify voice quality, and therefore we find LTAS as a suitable (and quite simple) method for detecting a decline in voice quality for severe OSA. Based on LTAS we extract the following features: the Absolute Spectral Slope (SLOPE.LTAS), statistical measures: spectral centroid (CENTRAL.LTAS), spectral spread (SPREAD.LTAS), spectral skewness (SKEWNESS.LTAS), spectral kurtosis (KURTOSIS.LTAS). Next, we have the spectral roll-off (ROLLOFF.LTAS) which, as the SLOPE.LTAS measure, quantifies the energy decay at higher frequencies. Finally, we have two measures computed on 5 frequency bands of the LTAS: the Spectral Flatness Ratio (SFR1–5) and Spectral Crest (SC1–5); the frequencies bands are: 175–500 Hz, 500–1000 Hz, 1000–2000 Hz, 2000–3000 Hz, and 3000–4000 Hz.

The nomenclature used for the features is as follows: for vowels we have *[measure].V[position].vowel* as in, for example, SFR_VE.O, while for the phrase we have *[measure].F[position]*, as in, for example, SC2.Itas.FA.

Classification problem

In order to quantify the utility of voice in detecting OSA, we focus primarily on the binary classification problem of the extreme groups: the control group and the severe OSA group. If voice were to be considered an important factor in detecting OSA, then it should discriminate well at least the most extreme categories.

Classifiers

The discrimination power is measured through experiments we perform with several classifiers.

The first classifier employed was a classical Multi Layer Perceptron (MLP) Neural Network trained with the Back Propagation technique with an adaptive learning rate [26]. As discussed in Feature reduction section we will perform a feature input-space reduction to 5 dimensions. We choose a two hidden layer MLP with $n_i:n_{h1}:n_{h2}:n_o$, where the number of inputs $n_i=5$, the number of nodes on the first hidden layer $n_{h1}=10$, the number of nodes on the second hidden layer $n_{h2}=5$, and the number of output nodes $n_o=2$ (as we have two classes). The activation function (transfer function) is the hyperbolic tangent sigmoid function. The inputs suffered a pre-normalization step, that is: all features values were linearly mapped to $[-1,+1]$; The MLP runs for $T_{MLP}=750$ iterations (epochs) found sufficient to achieve a low classification error margin and a good generalization for most of the runs.

Next, we apply a Support Vector Machine [27] classifier which is a powerful kernel-based classification paradigm. We used the simple linear kernel variant SVM (SVM_{lin}) that performs a linear discrimination, and the non-linear kernel variant (SVM_{poly}) which employs a polynomial kernel of degree 3, capable of finding non-linear decision boundaries between classes.

AdaBoost [28] is a classifier that combines several weak classifiers (in our implementation these weak classifiers are decision trees) to produce a powerful classification scheme with good generalization capabilities. AdaBoost is quite successful in modern face recognition applications [29].

We also employed a k-Nearest Neighbour (KNN) classification strategy [32] where the number of neighbours was taken to be 5.

Finally, we checked the performance of a classical Bayesian Classification (BC) scheme that uses a multivariate Gaussian model for the distribution of each class, assuming independence between features (a diagonal covariance matrix for the model, implying a linear decision boundary) [28,32].

Crossvalidation

In order to obtain a good estimate of the classifier's performance on a relatively reduced set of patterns, as the one employed in our study, we may first perform a crossvalidation process and then draw suitable conclusions on the mean classification errors obtained. We employ an S-fold crossvalidation method [28] that consists of dividing the ordered set of patterns into S contiguous chunks containing approximately the same number of patterns each, and then performing S training-testing experiments as follows: for each chunk $i \in \{1, 2, \dots, S\}$ we hold the current chunk for testing the classifier and we perform training on the remaining $S-1$ chunks, recording the results. We repeat the S training-testing experiments for a number of τ trials, each trial starting with a random permutation of the whole set of patterns. The main result of each training-test experiment is the Correct Classifications Rate (CCR) expressed as a percentage. The S-fold crossvalidation yields

a matrix of $\tau \times S$ of results from each training-testing experiment. We denote the matrix as **CCR**. The process is identical for all classifier but the MLP. It is well-known that neural networks are prone to get stuck in local minima of the error surface as basically they perform a gradient-descent or other similar local optimization with respect to the free parameters (weights, biases) [26]. Therefore, for a given set of training and test patterns it is important to perform several trials with different (usually random) starting points (values for weights and biases) and take into account the best run. In our case, the S-fold crossvalidation for the MLP performs ξ runs of the neural network with randomly taken starting points (random initialization), for each of the $\tau \times S$ training-testing experiment. After one such experiment we record only the best run in ξ runs. Thus, the matrix of recorded results **CCR** will still be $\tau \times S$ dimensional. For all experiments we took $S=5$, $\tau=50$, and $\xi=20$.

Feature reduction

Due to the high number of features employed in our study, which is 253, and the relatively low number of available subjects (248), in order to avoid the curse of dimensionality [30] (i.e. a uniform and sufficiently dense sampling in such high dimensional spaces, requires a huge number of data/patients), we must reduce the dimensionality of the feature space. We do so using two strategies: feature selection (find a small number of representative features) and feature combination (apply a transformation to the input feature space to produce a reduced output feature space). In all cases we perform a strong reduction from 253 to 5 variables (i.e. a 5-dimensional feature vector).

Feature ranking

The first method used to reduce the dimensionality of the feature space is a selection scheme that first ranks all features according to a statistical test of the discrimination power of each feature. Discrimination refers to the values each feature may take for the two classes involved in the comparison: control group and severe OSA. We observed that most of the features for both classes have a distribution that deviates significantly from the normal distribution and moreover they present outliers (Fig. 1a). Therefore, the test employed should not rely on normality assumptions, and we choose for a nonparametric test that is the two-sample unpaired Wilcoxon test (also known as the Mann–Whitney U-test) [31]. The method ranks the features in the entire set Φ of 253 features using the independent evaluation criterion for binary classification. This yields a number Z for each feature which is the absolute value of the u-statistic. Moreover, we outweigh the Z values using the following equation:

$$Z_{\text{final}} = Z \cdot (1 - \alpha\rho) \quad (1)$$

where $\alpha \in [0, 1]$ a parameter of the method and ρ is the Pearson cross-correlation coefficient between the candidate feature and all previously selected features. We took $\alpha=0.9$, that is we outweigh the significance statistic, meaning that features that are highly correlated with the features already picked are less likely to be included in the output list. Finally, we sort in decreasing order all features upon Z_{final} , taking the 5 features which correspond to the top 5 Z_{final} values.

Genetic algorithms-based feature selection

Genetic Algorithms (GAs) as part of the wider field of Evolutionary Algorithms (EAs) are population-based, stochastic search and optimization methods inspired by the natural evolution process [33]. The populations consist of a fixed number of potential

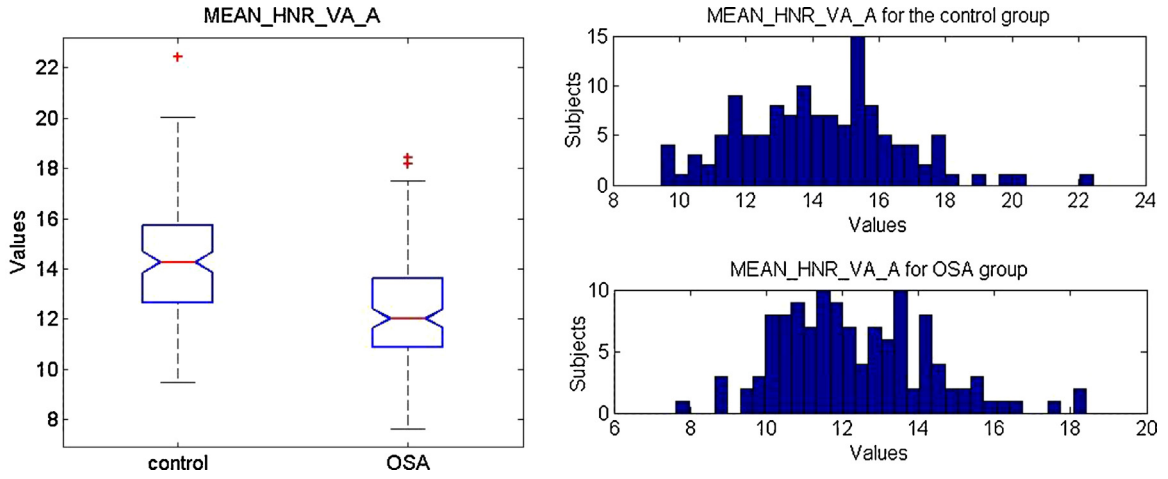


Fig. 1. Boxplot (a) and histograms (b) of the MEAN_HNR_VA_A features for the control and severe OSA group.

solutions to the optimization problem, called “chromosomes”. That is:

$$P = \{x_i \mid i = \overline{1 \dots N}\} \text{ and } x_i = (x_{i1}, \dots, x_{il}), x_{ij} \in [v_{lb_j}, v_{ub_j}] \subset \mathbf{R},$$

$$\forall i = \overline{1 \dots N}, \forall j = \overline{1 \dots l} \quad (2)$$

with N the size of the population P and x_i the chromosomes in P defined (for the present application) as vectors of integer genes x_{ij} ; v_{lb_j} and v_{ub_j} represent the lower and upper bound respectively of the genes' values. Each chromosome x_i bears a utility score $F(x_i)$ called *fitness* in direct relationship with the optimization criterion. It is expected that by repeated application of *selection* of the best chromosomes and variation operators called *crossover* and *mutation* to the whole population, the algorithm *evolves* such as the average fitness of the chromosomes increases/decreases (maximization/minimization). The final populations contain the optimal or near optimal solutions.

For feature selection purpose, each gene corresponds to the index of a feature in Φ , thus it is an integer between 1 and 253 (i.e. $v_{lb_j} = 1$, $v_{ub_j} = 253$, $\forall j$), and $l = 5$ as we want to reduce the dimension of the feature set to 5. That is, the GA seeks the best combination of 5 unique features from the entire set of available features Φ , according to an optimization criterion (fitness function). The termination criterion of the algorithm is the expiration of the maximum number of generations the GA is let to run (T_{\max}).

Selection is a probabilistic mechanism which chooses the best individuals (i.e. minimum fitness) with some probability from the current generation and passes them to the next generation. We have adopted a binary tournament selection scheme [33] due to its constant selection pressure over time [34]. We prevent losing the best individuals from the population [35] by an elitist replacement of the 5 worst individuals in each generation with the 5 best individuals in the previous generation. We used a rather high number of elites (i.e. 5) as we adopt a relative high mutation rate as well (see end of this subsection).

The *fitness function* we propose is in direct relationship with the classification performance. We choose the fitness function for a given chromosome x (i.e. a given combination of l features) to be proportional to the Error Rate (ER%) obtained after measuring how well a classifier discriminates the two classes using the features x . As we perform minimization we seek the best combination of l features that minimizes ER or equivalently minimizes the quantity $100 - \text{CCR}\%$. We evaluate the performance of the classifier by performing an S -fold crossvalidation as described in Crossvalidation

section. The fitness function should penalize the repetition of features in a chromosome x – we seek a vector of l *unique* features. It should also penalize a high variation of the CCR values in the S -fold crossvalidation for chromosome x , as we seek, besides high CCR values, a *reduced variation* of CCR between training/testing experiments in the crossvalidation. That is, $f(x)$ should increase substantially if we encounter repetitions of the features and should increase mildly with the variance of the CCR results after S -fold crossvalidation. The fitness (minimization) is taken as:

$$f(x) = 100 - \overline{\text{vec}(\text{CCR}(x))} + \underbrace{\omega \cdot \text{std}(\text{vec}(\text{CCR}(x)))}_{\text{penalty term1}} + \underbrace{\text{rep}(x) \cdot e^{\text{rep}(x)}}_{\text{penalty term2}}, \quad x \in P \quad (3)$$

where $f(x)$ is the fitness of feature vector x , “vec” represents the operator that stacks the matrix columns into a vector, the upper horizontal bar is the average operator, ω is the weight of the first penalty term which is the standard deviation of CCR, “rep” is the repetition operators that counts how many repeated features occur in the feature vector x , and is used for the second penalty term. By increasing $f(x)$ through the penalty terms, due to the selection effect in the population, such “bad” chromosomes tend to disappear after several generations of the GA.

The variation operators are: Uniform Crossover (AX) defined in [36] and applied to pairs of chromosomes with some probability P_c . Mutation (flip mutation) [33] replaces, with some probability P_m , the gene's value at a given locus j with a random value in $[v_{lb_j}, v_{ub_j}]$.

The GA population is initialized as follows: half of the chromosomes in the population, chosen at random, get their initial gene's values by picking randomly for each gene the index of a feature in the top 62 (approximately a quarter of the features in Φ) of previously ranked features. Features are ranked according to the decreasing ordered set of Z_{final} values as described in Feature ranking section. Thus, we assure that at least half of the population has been initialized with good features, the rest of the population being initialized with random values in $[v_{lb_j}, v_{ub_j}]$. Random initialization is a standard procedure for GA that allows for a wide exploration of the search space from the first generation.

The parameterization of the GA in all our experiments is the following: $N = 50$ individuals, $T_{\max} = 100$ generations, $P_c = 0.9$, $P_m = 0.2$,

$\omega=0.5$; the S -fold crossvalidations when calculating the fitness have $S=5$, $\tau=10$, and in the case of the neural network (MLP) $\xi=5$.

Feature combination

We may reduce the dimensionality of the feature space by performing linear transformation and taking the most important components. We adopt Principal Component Analysis (PCA) which is a well-known statistical technique that has been widely used in data analysis and compression (for example, articles such as [37] and textbooks such as [28] present reviews of the method). The goal of the method is the compression of a high-dimensional input data into a lower dimensional space, without loss of relevant information. To capture the main features of the data set, PCA is looking for directions along which the dispersion or variance of the point cloud is maximal. These “principal” directions form a subspace of lower dimension than the original input space. The projection of the data onto the respective subspace will yield a transformation similar to compression, which minimizes the loss of information according to the Minimum Mean Square Error criterion. In our case, we perform the transformation over the 253-dimensional feature space and take only the first 5 principal components.

Fisher’s Linear Discriminant Analysis (LDA) is a well known dimensionality reduction scheme [32] that projects the patterns onto a lower dimensional subspace such that the classes become “more separable” according to a criterion (maximization) called the Fisher Linear Discriminant.

Results

Discriminating potential of voice features

We may analyze the discriminating potential of voice features (Voice features section) by looking at the top 5 features as yielded by the ranking method described in Feature ranking section. These features (in decreasing order of the Z_{final} values) are the following: ‘MEAN.HNR.VA.A’, ‘VTI.VE.A’, ‘MBW.formant2.VA.E’, ‘MBW.formant2.VE.I’, ‘MF.formant2.VE.U’. The best feature is therefore ‘MEAN.HNR.VA.A’ that passes (favours the alternative hypothesis) the Wilcoxon two-sampled test of difference in medians with a good p -value, $p=2.09 \times 10^{-10}$ (the null hypothesis states that medians are equal for the two groups – control and severe OSA – and we reject the null hypothesis at a 1% significance level with a quite small p -value, $p<0.01$). From the boxplot in Fig. 1, the difference between the distribution of ‘MEAN.HNR.VA.A’ is apparent. Moreover, from histograms in Fig. 1a and b, it is apparent that distributions for the two groups seem to depart from normality and outliers occur. This justifies the use of a non-parametric robust statistical test, in the first place (i.e. Wilcoxon test). Furthermore, besides the mean and standard deviation values we gave the median values as well, less affected by outliers and heavy tail skewed distributions. It is relevant to note that the use of ‘MEAN.HNR.VA.A’ was inspired by studies that try to discriminate between normal voice and sleepy voice [38], as we consider that severe OSA patients may exhibit certain fatigue in voice. Looking at the next 4 features and applying the same statistical test, it follows that for all features, the medians between groups cannot be considered equal, and this is a strong assertion judging by the very small p -values obtained ($p<0.00001$): 2.057×10^{-12} (‘VTI.VE.A’), 1.4843×10^{-8} (‘MBW.formant2.VA.E’), 2.77×10^{-8} (‘MBW.formant2.VE.I’), 1.9×10^{-6} (‘MF.formant2.VE.U’). Even though ‘VTI.VE.A’ has a smaller p -value than ‘MEAN.HNR.VA.A’, the ranking method outweighed this feature as it was found to be correlated to ‘MEAN.HNR.VA.A’. Statistics performed indicate that, at least taking into account the top 5 ranked features, voice may be

Table 1

Considered database for this study. Gender, Age (range, mean, median and mode) and Body Mass Index (range, mean, median and mode) of both groups are provided.

Control group	
Gender	48 men, 79 women
Ages	18 ÷ 64, mean 29.68, median 24, mode 21
Body Mass Index (BMI)	18 ÷ 64, mean 29.68, median 24, mode 21
Severe OSA group	
Gender	101 men, 20 women
Ages	28 ÷ 82, mean 54.04, median 55, mode 62
Body Mass Index (BMI)	23 ÷ 53, mean 32.56, median 31.2, mode 34.6

considered *distinct* between the extreme groups: control and severe OSA (Table 1).

Classifier comparison

The results of the S -fold crossvalidation for all classifiers and feature reduction schemes are given in Table 2. It follows that the best strategy in terms of CCR (average 82.04%), Sensitivity (average 81.74%) and Specificity (average 82.40%) is the Bayesian Classifier (BC) with featured selected by the GA (denoted as BC-GA for short). BC-GA also achieves the smallest standard deviations of the CCR/Sensitivity/Specificity triplet among all classifiers and

Table 2

Results in terms of Average (AVG), Median (MED), Mode (MOD), Standard Deviation (STD) for the Correct Classification Rate (CCR), Sensitivity and Specificity for all classifiers and feature reduction methods.

		Ranked	PCA	LDA	GA
MLP	CCR [%]	AVG	78.98	77.37	79.98
		MED	79.59	77.77	80
		MOD	76	78	82
		STD	4.64	4.95	5.03
	Sensitivity [%]	AVG	75.01	73.77	83.75
		MED	77.27	76.92	86.36
		MOD	83.33	80	87.5
		STD	13.17	14.56	10.4
	Specificity [%]	AVG	75.58	76.64	78.3
		MED	77.27	78.26	89.32
		MOD	81	76.92	81
		STD	11.46	13.82	10.59
	SVM _{lin} CCR [%]	AVG	74.22	72.66	77.16
		MED	74	73.46	77.55
		MOD	76	74	78
		STD	6.12	5.02	5.38
	Sensitivity [%]	AVG	72.29	71.97	74.9
		MED	72.72	72	75
		MOD	75	66.66	75
		STD	9.6	8.56	8.73
	Specificity [%]	AVG	76.34	73.55	79.14
		MED	76.92	74	79.31
		MOD	78.26	74	81
		STD	8.87	9.45	6.91
SVM _{poly}	CCR [%]	AVG	65.35	68.74	74.05
		MED	66	69.38	74
		MOD	66	70	74
		STD	6.67	6.12	5.54
	Sensitivity [%]	AVG	63.76	70.14	71.82
		MED	65	70.83	71.19
		MOD	66.66	70	71
		STD	10.51	9.18	8.92
	Specificity [%]	AVG	67.06	67.65	76.44
		MED	66.66	68.18	77.77
		MOD	66.66	66.66	78.57
		STD	10.46	9.07	8.82

Table 2 (Continued)

		Ranked	PCA	LDA	GA	
AdaBoost	CCR [%]	AVG	72.42	71.61	76.28	
		MED	72	71.42	76	
		MOD	74	74	80	
		STD	5.87	5.94	5.56	
Sensitivity [%]	AVG	68.24	70.48	74.18	76.18	
		MED	69.56	71.42	75	76.66
		MOD	75	75	75	76
		STD	9.92	9.83	8.86	8.95
Specificity [%]	AVG	76.8	70.22	78.43	78.86	
		MED	76.92	71.13	78.57	79.16
		MOD	75	72	75	80
		STD	9.13	9.43	8.21	8.7
KNN	CCR [%]	AVG	62.58	67.98	74.62	
		MED	63.26	67.34	74	
		MOD	66	66	75	
		STD	5.89	5.72	5.76	
Sensitivity [%]	AVG	63.01	70.6	72.12	70.72	
		MED	62.5	70.37	72	70.83
		MOD	66.66	70	72	72
		STD	10.40	9.46	9.45	9.20
Specificity [%]	AVG	62.78	65.96	77.42	78.03	
		MED	62.06	65.45	77.95	77.77
		MOD	60	60	78	80
		STD	8.72	9.25	8.52	8.06
BC	CCR [%]	AVG	71.71	73.69	77.57	
		MED	72	73.47	78	
		MOD	70	73	78	
		STD	5.72	5.31	5.43	
Sensitivity [%]	AVG	75.63	76.54	75.17	81.74	
		MED	76.19	77.27	76	82.60
		MOD	75	77	75	82.60
		STD	8.37	8.22	9.28	7.18
Specificity [%]	AVG	67.96	71.03	79.76	82.40	
		MED	68	71.19	80	82.60
		MOD	67	72	80.76	87.50
		STD	8.02	8.50	7.45	7.57

all feature selection methods. The second best strategy is the SVM with linear kernel (SVM_{lin}) and features selected by the GA, and the third best is the MLP with a dimensionality reduction through LDA. GA achieves the best or close to best 5 features for each classifier, and therefore is the best feature selection scheme, while LDA is the best feature combination method.


A closer look at the results of the BC-GA (Fig. 2) indicates that Sensitivity (Fig. 2a) and Specificity (Fig. 2b) present a skewed distribution with longer tails for smaller than 70% values (the skewness is -0.47 for Sensitivity and -0.21 for Specificity), therefore we might consider that the most representative (probable) values for the Sensitivity and Specificity are closer to the mode than to the mean of the distribution (Sensitivity $\approx 83\%$, and Specificity $\approx 88\%$).

It is instructive to see what are the 5 features selected by the GA in the case of BC-GA: 'CVF_pitch_VE.A', 'VTI_VE.A', 'MBW_formant2_VE.E', 'APQ_VE.I', 'SC2.Itas.FA'. One of the features is related to the phrase analysis and is the Spectral Crest of the LTAS for the second frequency band 500–1000 Hz ('SC2.Itas.FA', see LTAS based section). Thus, both vowel processing and phrase processing are important in decision making.


The best triplet (CCR, Sensitivity and Specificity) in a single training-test run of the BC-GA during the S-fold crossvalidation is CCR = 96%, Sensitivity = 96% and Specificity = 96%.

Next, we consider an ensemble of the three best classifiers: BC-GA, SVM_{lin}-GA and MLP-LDA. We apply majority vote for the outputs of the classifiers to decide the class to which each pattern

Crossvalidation results (CCR) for the Bayes Classifier with GA feature reduction



Crossvalidation results (Sensitivity) for the Bayes Classifier with GA feature reduction



Crossvalidation results (Specificity) for the Bayes Classifier with GA feature reduction




Fig. 2. Histograms of the (a) CCR, (b) Sensitivity and (c) Specificity for the S-fold cross-validation of the Bayesian Classifier with features selected by the GA.

pertains. For the whole ensemble we apply the S-crossvalidation process as before and get a slight improvement of the results as given in Table 3. The best triplet (CCR, Sensitivity and Specificity) in a single training-test run of the ensemble classifier during the S-fold crossvalidation is CCR = 95.91%, Sensitivity = 92.85% and Specificity = 100%.

Next, we check the performance on “in-between cases”, that are snoring patients and mild OSA. These cases have an AHI between 5 and 30. We check the performance of the BC-GA on a group of 128 patients, where we consider AHI = 15 as the border between non-OSA (below this threshold) and OSA (above this threshold). We have 65 non-OSA and 63 OSA patients, 24 women and 104 men, with ages comprised between 21 and 76, mean 48, median 48 and mode 42. BMI ranges from 18 to 47, mean 29, median 28 and mode 26. We train the BC-GA on all 248 extreme cases (patterns considered in previous sections) and test the classifier on the 128 intermediate cases only.

We get a CCR = 70.31%, with a Sensitivity = 73.01% and a Specificity = 67.69%. This is a clear drop in performance with respect to validation on extreme cases, meaning that it is difficult to

Table 3

Results in terms of Average (AVG), Median (MED), Mode (MOD), Standard Deviation (STD) for the Correct Classification Rate (CCR), Sensitivity and Specificity for the ensemble classifier (BC-GA + SVM_{lin}-GA + MLP-LDA).

		Ensemble
CCR [%]	AVG	82.85
	MED	82
	MOD	83
	STD	4.83
Sensitivity [%]	AVG	81.49
	MED	81.48
	MOD	84
	STD	7.57
Specificity [%]	AVG	84.69
	MED	85.71
	MOD	87
	STD	6.52

Table 4

Results in terms of Average (AVG), Median (MED), Mode (MOD), Standard Deviation (STD) for the Correct Classification Rate (CCR), Sensitivity and Specificity for the BC-GA classifier on the 128 intermediate cases.

		BC-GA
CCR [%]	AVG	64.23
	MED	64.69
	MOD	69.23
	STD	8.26
Sensitivity [%]	AVG	55.36
	MED	55.55
	MOD	60
	STD	14.34
Specificity [%]	AVG	72.94
	MED	73.33
	MOD	69.23
	STD	11.68

discriminate between milder OSA and non-OSA/snoring patients based solely on knowledge acquired from the voice in the extreme cases group. Moreover, we may perform an *S*-fold crossvalidation ($S=5$, $\tau=50$) on the intermediate patients alone. In this case, we get the results presented in Table 4, which again show a dramatic drop in performance. We may conclude at this point that it is hard to build an efficient classifier using the intermediate cases alone, and is preferable to build the classifier on the extreme cases (this assures, at least, a good recognition rate for the extreme cases: CCR above 80%) and a recognition rate (CCR) for the intermediate validation cases around 70%. For completion, we also perform a *S*-fold crossvalidation on *all* 376 patterns (union between extreme and intermediate cases). Results are given in Table 5.

Table 5

Results in terms of Average (AVG), Median (MED), Mode (MOD), Standard Deviation (STD) for the Correct Classification Rate (CCR), Sensitivity and Specificity for the BC-GA classifier on the 128 intermediate cases +248 extreme cases.

		BC-GA
CCR [%]	AVG	74.9
	MED	74.66
	MOD	73.33
	STD	4.92
Sensitivity [%]	AVG	71.39
	MED	71.42
	MOD	70
	STD	8.34
Specificity [%]	AVG	78.23
	MED	78.57
	MOD	80
	STD	5.99

The classifier trained on extreme cases achieves the best results when validated on extreme cases, and significantly worse results when validated on intermediate cases. A potential means of sieving-out the intermediate cases prior to application of the classifier (in a screening scenario, for example) would be the use of simple parameters from the patient's medical or clinical record, readily available standard measures such as age, BMI, neck circumference, Epworth sleepiness scale (EPW), blood pressure, etc. Numerous studies such as [39–43] are investigating the relationships and correlations between such standard measures and the OSA. We may perform a quick check of these relationships by computing the Pearson linear correlation coefficient (ρ) between basic measures and AHI on our body of 248 extreme cases. We obtain $\rho=0.60$ (p -value $<10^{-25}$) for the correlation between AHI and age, $\rho=0.49$ (p -value $<10^{-20}$), for the correlation between AHI and BMI, and $\rho=0.29$ (p -value $<10^{-20}$), for the correlation between AHI and EPW. These values indicate that we can rely on the assumptions that these measures (especially age) are correlated to AHI, therefore they may be employed to cull-out potential intermediate cases. For example, cases with ages between 30 and 45, relatively low or medium BMI, medium EPW values, evidence of snoring during nightsleep, may be discarded as intermediate cases prior to the application of the classifier. Such intermediate cases need a deeper analysis, and fusion with other more traditional sources of information (such as the established means of diagnosis: RP, or PSG).

Finally, it is worth mentioning that more complex classifiers such as SVM_{poly} or KNN that attain nonlinear decision boundaries have less generalization capabilities than simpler classifiers such as BC or SVM_{lin} that use linear decision boundaries. The explanation is that such complex classifiers seem to present *overfitting*, in that they are capable of learning very well the training patterns with all incorporated noise and spurious information, but the complex decision border is not able to classify well the new test patterns. Simpler (linear) decision borders seem better for the current distribution of patterns. Actually, by looking at several runs in the *S*-fold crossvalidation of the MLP, we found that many times overfitting of the neural networks [44] occurs: when monitoring the learning curve, at some point in time as the learning error keeps decreasing the classification error of the test patterns starts to increase. For such runs we performed early-stopping, that is we stopped learning for an epoch less than T_{MLP} when the error on the test set began increasing.

Conclusions

The present study focuses on voice alone as a primary discriminating source of information between healthy subjects and severe OSA. Both statistical analysis on several voice extracted features, as well as performance of several classifiers indicate that voice has a clear potential to detect severe OSA among healthy subjects. The performance of the classifiers has been estimated using robust statistical techniques (*S*-fold crossvalidation) while counting with a relatively large body of subjects (i.e. 248), larger than most of the present studies analyzing the relationship between voice and OSA. The group of subjects involved in our experimental design increases to 376, when including the intermediate cases as well. We may get a better grasp on the relationship between OSA and voice by looking at the extreme cases that also have a clear-cut diagnosis. The results in terms of CCR, Sensitivity and Specificity, all above 80% for several classifiers point out the good potential of voice as a discriminating factor between healthy subjects and severe OSA.

Careful analysis on subjects with different degrees of OSA reinforced our prior belief that voice may act as a good discriminating factor for most of the severe cases. However, for intermediate cases where upper-airway closure may not be so pronounced (thus voice

not much affected), we cannot rely on voice alone for making a good discrimination between OSA and non-OSA.

Analysing the features discovered by the feature reduction methods, we conclude that both vowel and phrase features are useful (more vowel features are selected, however) and both uttering positions as well, with more features selected from the stretched ('E') uttering.

The GA feature selection method proved to be the best reduction scheme that is well adapted to the classifier, and that achieves the best CCR, Sensitivity and Specificity with a small variance of these results due to the specifically designed fitness function (see Eq. (3)), for almost all cases involved in comparison. The GA is capable of discovering useful associations between voice features, and that are not apparent beforehand, the degree of utility being in direct relationship to the classifier performance.

Feature selection is a crucial stage in our design as there are many features that can be extracted from voice and speech but there is no *a priori* knowledge regarding the most discriminant to be employed in the detection of the OSA cases. Therefore, we would like to highlight the use of GAs as one of the most innovative aspects in the present study. GAs have turned out to be the perfect choice when it comes to salient feature discovery, achieving good adaptation with the classification tools employed.

For a screening application that detects severe OSA cases among healthy people we may employ an ensemble classifier that combines the output of various classifiers to yield a more robust decision. As seen from Classifier comparison section such an ensemble classifier achieves slightly better results than the best classifier (BC-GA). Moreover, fusion with other measures from the subject's medical record (i.e. sex, age, BMI, EPW, blood pressure) is expected to increase the overall performance. Such parameters are correlated with the AHI index and thus with the presence or absence of OSA, and may shed light into the suitable discrimination of the intermediate subjects as well (mild OSA, snoring subjects), subjects that are difficult to classify by voice analysis only. A multiclass approach, instead of a binary classification, is also expected to increase the classification performance. We might consider more than 2 classes, such as, for example: controls, healthy snoring subjects, mild-OSA, and severe-OSA, and we may make a differentiation between sexes, as well.

So far, results presented as an S-fold crossvalidation for several classifiers are by no means a substitute for a clinical validation study. Crossvalidation served us to better estimate the discriminating potential of voice, and the *expected* correct classification rate, sensitivity and specificity. Actually, during each training-testing experiment involved in the S-fold crossvalidation only a fifth of the total number of subjects (about 50, for the extreme cases problem) was employed for validation purposes, the rest being used to train the classifier. For future work, we will seek to produce clinical validation results for a comprehensive body of new subjects, with an already trained classifier using the model developed in this paper.

References

- [1] S.M. Caples, A.S. Gami, V.K. Somers, Obstructive sleep apnea, *Ann. Intern. Med.* 142 (3) (2005) 187–197.
- [2] C.A. Kushida, M.R. Littner, T. Morgenthaler, C.A. Alessi, D. Bailey, J. Coleman, L. Friedman, M. Hirshkowitz, S. Kapen, M. Kramer, T. Lee-Chiong, D.L. Loubé, J. Owens, J.P. Pancer, M. Wise, Practice parameters for the indications for polysomnography and related procedures: an update for 2005, *Sleep* 28 (4) (2005) 499–521.
- [3] P.K. Monoson, A.W. Fox, Preliminary observation of speech disorder in obstructive and mixed sleep apnea, *Chest* 92 (1987) 670–675.
- [4] A.W. Fox, P.K. Monoson, C.D. Morgan, Speech dysfunction of obstructive sleep apnea, *Chest* 96 (1989) 589–595.
- [5] J.A. Fiz, J. Morera, J. Ahad, A. Belsunces, M. Ham, J.I. Fiz, R. Jane, P.M. Caminal, D. Rodenstein, Acoustic analysis of vowel emission in obstructive sleep apnea, *Chest* 104 (1993) 1093–1096.
- [6] L.B. José, A.H. Luis, F. Rubén, R. Daniel, Improving automatic detection of obstructive sleep apnea through nonlinear analysis of sustained speech, *Cognit. Comput.* 5 (December (4)) (2013) 458–472.
- [7] Goldshtein E1, A. Tarasiuk, Y. Zigel, Automatic detection of obstructive sleep apnea using speech signals, *IEEE Trans. Biomed. Eng.* 58 (May (5)) (2011) 1373–1382.
- [8] A.D. Lucey, A.J.C. King, G.A. Tetlow, J. Wang, J.J. Armstrong, M.S. Leigh, A. Paduch, J.H. Walsh, D.D. Sampson, P.R. Eastwood, D.R. Hillman, Measurement, reconstruction, and flow-field computation of the human pharynx with application to sleep apnea, *IEEE Trans. Biomed. Eng.* 57 (10) (2010) 2535–2548.
- [9] Sleep-related breathing disorders in adults: recommendations for syndrome definition and measurement techniques in clinical research. The Report of an American Academy of Sleep Medicine Task Force *Sleep* 22(5) (1999) 667–589.
- [10] F.J. Tolkmitt, K.R. Scherer, Effect on experimentally induced stress on vocal parameters, *J. Exp. Psychol. Hum. Percept. Perform.* 12 (3) (1986) 302–313.
- [11] T. Kitamura, H. Takemoto, K. Honda, Y. Shimada, I. Fujimoto, Y. Syakudo, S. Masaki, K. Kuroda, N. Oku-uchi, M. Senda, Difference in vocal tract shape between upright and supine postures: observations by an open-type MRI scanner, *Acoust. Sci. Technol.* 26 (5) (2005) 465–468.
- [12] S.I. Buchaillard, P. Perrier, Y. Payan, A biomechanical model of cardinal vowel production: muscle activations and the impact of gravity on tongue positioning, *J. Acoust. Soc. Am.* 126 (4) (2009) 2033–2051.
- [13] L.R. Rabiner, R.W. Schafer, *Digital Processing of Speech Signals*, Prentice Hall, 1978.
- [14] D.D. Deliyski, Acoustic model and evaluation of pathological voice production, in: *Proceedings of EUROSPEECH'93*, Berlin, Germany, 1993, pp. 1969–1972.
- [15] J. Koreman, M. Pützer, Finding correlates of vocal fold adduction deficiencies, in: *PHONUS 3*, Institute of Phonetics, University of the Saarland, Saarbrücken, 1997, pp. 155–178.
- [16] D. Talkin, A Robust Algorithm for Pitch Tracking (RAPT), in: W.B. Kleijn, K.K. Paliwal (Eds.), *Speech Coding & Synthesis*, Elsevier, Amsterdam, 1995, pp. 495–518.
- [17] R.J. Bakken, R.F. Orlikoff, *Clinical Measurement of Speech and Voice*, 2nd ed., Singular, Thomson Learning, San Diego, CA, USA, 1999.
- [18] Y. Qi, R.E. Hillman, Temporal and spectral estimations of harmonics-to-noise ratio in human voice signals, *J. Acoust. Soc. Am.* 102 (1) (1997) 537–543.
- [19] H. Muta, T. Baer, K. Wagatsuma, T. Muraoka, H. Fukuda, A pitch-synchronous analysis of hoarseness in running speech, *J. Acoust. Soc. Am.* 84 (4) (1988) 1292–1301.
- [20] G. de Krom, A cepstrum-based techniques for determining a harmonics-to-noise ratio in speech signals, *J. Speech Hear. Res.* 36 (1993) 254–266.
- [21] V. Parsa, D.J. Jamieson, Acoustic discrimination of pathological voice: sustained vowels versus continuous speech, *J. Speech Lang. Hear. Res.* 44 (2001) 327–339.
- [22] M.A. Little, P.E. McSharry, S.J. Roberts, D.A.E. Costello, I.M. Moroz, Exploiting nonlinear recurrence and fractal scaling properties for voice disorder detection, *BioMed Eng. OnLine* 6 (23) (2007), Available from <http://www.biomedical-engineering-online.com/content/6/1/23>
- [23] T. Leino, Long-term average spectrum in screening of voice quality in speech: untrained male university students, *J. Voice* 23 (6) (2009) 671–676.
- [24] P. Boersma, G. Kovacic, Spectral characteristics of three styles of Croatian folk singing, *J. Acoust. Soc. Am.* 119 (3) (2005) 1805–1816.
- [25] G. Peeters, A Large Set of Audio Features for Sound Description (Similarity and Classification) in the CUIDADO Project, IRCAM, Tech Rep, 2004.
- [26] S. Haykin, *Neural Networks: A Comprehensive Foundation*, 2nd ed., Prentice Hall, Upper Saddle River, NJ, USA, 1999, pp. 178–277.
- [27] N. Cristianini, J. Shawe-Taylor, *An Introduction to Support Vector Machines and Other Kernel-based Learning Methods*, Cambridge University Press, Trumpington Street, Cambridge, UK, 2000.
- [28] C.M. Bishop, *Pattern Recognition and Machine Learning*, Springer-Science + Business Media, LLC, Secaucus, NJ, USA, 2006.
- [29] P. Viola, M.J. Jones, Robust real-time face detection, *Int. J. Comput. Vision* 57 (2) (2004) 137–154.
- [30] S. Theodoridis, K. Koutroumbas, *Pattern Recognition*, 2nd ed., Elsevier-Academic Press, Amsterdam, 2003, pp. 43.
- [31] P.H. Kvam, B. Vidakovic, *Nonparametric Statistics with Applications to Science and Engineering*, Wiley-Interscience, A John Wiley & Sons, Inc., Publ, Hoboken, NJ, 2007, pp. 129–133.
- [32] R.O. Duda, P.E. Hart, D.G. Stork, *Pattern Classification*, 2nd ed., Wiley-Interscience, New York, NY, 2001.
- [33] T. Bäck, *Evolutionary Algorithms in Theory and Practice*, Oxford University Press, Oxford, UK, 1996.
- [34] T. Bäck, L. Thiele, A comparison of selection schemes used in evolutionary algorithms, *Evol. Comput.* 4 (4) (1996) 361–394.
- [35] T. Bäck, F. Hoffmeister, Extended selection mechanisms in genetic algorithms, in: J.D. Schaffer (Ed.), *Proceedings of the 4th International Conference on Genetic Algorithms*, Morgan Kaufmann, San Diego, CA, 1991, pp. 92–99.
- [36] G. Syswerda, Uniform crossover in genetic algorithms, in: R.K. Belew, L.B. Booker (Eds.), *Proceedings of the Third International Conference on Genetic Algorithms*, Morgan Kaufmann, Fairfax, VA, 1989, pp. 2–9.
- [37] J.C. Bezdek, N.R. Pal, A note on self-organizing semantic maps, *IEEE Trans. Neural Networks* 6 (5) (1995) 1029–1036.
- [38] J. Krajewski, R. Wieland, A. Batliner, An acoustic framework for detecting fatigue in speech based human-computer interaction, in: K. Miesenberger, J. Klaus, W. Zagler, A. Karshmer (Eds.), *Computers Helping People with Special Needs*, Springer, Heidelberg, 2008, pp. 54–61.

- [39] E.O. Bixler, A.N. Vgontzas, T. Ten Have, K. Tyson, A. Kales, Effects of age on sleep apnea in men: I. Prevalence and severity, *Am. J. Respir. Crit. Care Med.* 157 (1) (1998) 144–148.
- [40] J.C. Ware, R.H. McBryer, J.A. Scott, Influence of sex and age on duration and frequency of sleep apnea events, *Sleep* 23 (2) (2000) 165–170.
- [41] I.L. Mortimore, I. Marshall, P.K. Wraith, R.J. Sellar, N.J. Douglas, Neck and total body fat deposition in nonobese and obese patients with sleep apnea compared with that in control subjects, *Am. J. Respir. Crit. Care Med.* 157 (1) (1998) 280–283.
- [42] T. Young, P.E. Peppard, D.J. Gottlieb, Epidemiology of obstructive sleep apnea: a population health perspective, *Am. J. Respir. Crit. Care Med.* 165 (9) (2002) 1217–1239.
- [43] K.P. Pang, D.J. Terris, Screening for obstructive sleep apnea: an evidence-based analysis, *Am. J. Otolaryngol.* 27 (2) (2006) 112–118.
- [44] C.M. Bishop, *Neural Networks for Pattern Recognition*, Oxford University Press, Oxford, UK, 1995, pp. 11–14.